Finding undetected protein associations in cell signaling by belief propagation

M. Bailly-Bechet [*], **C. Borgs** [†], **A. Braunstein** [‡], **J. Chayes** [†], **A. Dagkessamanskaia** [§], **J.-M. François** [§] **and R. Zecchina** [¶]

[*]LBBE, CNRS UMR 5558, Université Lyon 1, Villeurbanne, France,[†]Microsoft Research New England, One Memorial Drive, Cambridge, USA,[‡]Human Genetics Foundation, Via Nizza 52, 10126 Torino, Italy & Politecnico di Torino, C.so Duca degli Abruzzi 24, Torino, Italy,[§]Université of Toulouse, UMR CNRS-INSA 5504 & INRA 792, Toulouse, France, and [¶]Politecnico di Torino, C.so Duca degli Abruzzi 24, Torino, Italy & Human Genetics Foundation, Via Nizza 52, 10126 Torino, Italy

External information propagates in the cell mainly through signaling cascades and transcriptional activation, allowing it to react to a wide spectrum of environmental changes. High-throughput experiments identify numerous molecular components of such cascades which however may interact through other undetected molecules. Some of these interactions may be detected using data coming from the integration of a protein-protein interaction network and mRNA expression profiles. This inference problem can be mapped onto the one of finding appropriate optimal connected sub-graphs of a network defined by these datasets. The optimization procedure turns out to be computationally intractable in general. Here we present a new distributed algorithm for this task, inspired from statistical physics, and apply this scheme to alpha factor and drug perturbations data in yeast. We identify the role of the COS8 protein, a member of a gene family of previously unknown function, and validate the results by genetic experiments.

The algorithm we present is specially suited for very large data sets, can run in parallel and can be adapted to other problems in systems biology. On renowned benchmarks it outperforms other algorithms in the field.

COMPUTATIONAL BIOLOGY | BELIEF PROPAGATION | STEINER TREES | PHEROMONE PATHWAY

Introduction

Signaling cascades, an exemplar of which is the phosphorylation MAPK kinase pathways, consist in sequential reactions starting at receptor proteins and transmitted through protein interactions to effector proteins. Activation of these effectors leads to cellular changes, notably at the transcriptional level, and results in the adaptation of the cell to its surroundings [1]. Identifying signaling pathways is particularly important for medical studies since their malfunction is responsible for many diseases, such as cancer [2], or Alzheimer [3].

From an engineering point of view, signaling cascades present interesting properties: they provide signal filtering [4] and amplification [5]. Their global interconnected organization equip the cell with an integrated sensor network where pathways can modulate one another through cross-talk and retroactions. In this complex system, signal specificity is maintained by scaffold proteins [6, 7] acting as connectors of particular reactions. Finally, the output of the information carried by the transduction network allows for another layer of regulation, namely combinatorial control in gene expression [8]. A purely experimental approach to the identification of all components of a pathway, or all components of a functional gene module, would need long and costly experiments. Such process would greatly benefit from the extraction of indirect information about pathways from producible large-scale data, such as expression, sequencing and protein interaction data. Indeed, even if the correlation between signaling pathway activity and expression data may be weak (although this may be case-dependent, see [9] as an interesting example), expression data are available in shear quantity. By introducing a parameter weighting the relative importance given to expression data with respect to reliable protein-protein interaction data or in general established pathways knowledge, we hope to be able to enrich the existing knowledge by the small amount of information needed to reveal unknown interactions. To this scope, important aspects like the varying reliability of interaction data and the proliferation of alternative paths which need to be compared, require the development of heuristic algorithmic techniques which need to be efficient on large scale data sets.

Here we propose a new method for the inference of hidden components of functional networks and signaling pathways, from large-scale transcriptomics and protein interaction data. Such functional networks, composed of proteins acting together in given environmental conditions, are an integrated way of describing information processes in the cell. This problem has enormous potential applications and has already been addressed in several works [10, 11, 12, 13], leading to interesting theoretical predictions. In these works, the underlying methodology consists in separately identifying single signaling pathways, and then collecting them in an aggregated network. The methodology proposed here attempts instead to extract information on an entire network, defined as a connected sub-graph of the full protein interaction network. This technique was roughly sketched in a biological inference context[14]. Here we present a complete, in-depth description and analysis of the approach, including in particular algorithmic and experimental validations, and comparison with results from a previous work along the same lines. We also give full details of the algorithmic framework we use, which may allow implementation of the same ideas to similar systems biology problems.

We state the functional network inference problem in a rather simple and general graphical form. Given a graph $G = (V, E)$ – the protein interaction network (PIN) – with positive costs over edges $\{c_e : e \in E\}$ and positive prizes over vertices $\{b_i : i \in V\}$, find a connected sub-graph $G' = (V', E')$ that minimizes the following function:

$$\min_{\substack{E' \subseteq E, V' \subseteq V \\ (E', V') \text{ connected}}} \sum_{e \in E'} c_e - \lambda \sum_{i \in V'} b_i \qquad [\mathbf{1}]$$


Reserved for Publication Footnotes

To our purpose the costs of edges c_e are chosen so that high confidence interactions (protein interactions verified in small-scale experiments or found in many large-scale datasets) have lower value with respect to low confidence ones (interactions experimentally shown only once in a large-scale experiment). The node prizes are computed by $b_i = -\log p_i$, where p_i is the p-value of differential expression of node i in the corresponding microarray. The parameter λ regulates the tradeoff between the edge costs and vertices prizes, and its value indirectly controls the size of the sub-graph G'.

In spite of its apparent simplicity, the problem of solving Eq. [1], known as *Prize-collecting Steiner Tree* problem (PCST), is computationally intractable (NP-Hard) and heuristic algorithms need to be developed in order to solve instances arising from large datasets. Satisfying the connectivity constraint on the optimization task constitutes a major computational difficulty. The problem remains intractable even in the case in which $b_i \in \{0, L\}$ for a large $L > 0$, as this limit case corresponds to the better known *Minimum Steiner Tree* problem (MSTT) on graphs, which is also NP-Hard.

Modeling ideas related to our work are discussed in [15, 16, 17, 18]. These studies rely on different algorithmic techniques, namely on a combination of linear programming relaxation solvers, branch and bounds optimization methods and preprocessing of the underlying biological network. A detailed comparison shows essentially that the difference in performance between LP-based methods and the one presented here increases dramatically with the problem size. Additionally, the computational cost of our approach scales roughly linearly with the size of the problem and the algorithm is fully parallelizable. These two facts suggest that the method proposed here may be particularly well-suited to study problems defined on large networks.

The paper is organized as follows: First we present the general problem of identifying optimal subgraphs as a technique for integrating different data types. Then we discuss a new algorithmic approach based on belief propagation: we provide benchmark performances together with a specific application to pheromone response data in yeast. Finally we describe in detail the experimental validation of the predictions relative to the functional role of a family of genes (COS). Complete details are given in the Supporting Information.

Fig. 1. An example of a prize-collecting Steiner tree. Larger nodes mean larger prizes, thickness of the edges is proportional to their cost. A prize collecting Minimum Steiner Tree (right) picks as many as possible of the larger nodes while simultaneously picking the thinnest links and maintaining connectivity. The analyzed yeast protein network has approximately 5000 nodes and 22,000 edges

Results

A Message-Passing algorithm for PCST. In [19], a statistical physics analysis of the properties of Steiner Trees on different ensembles of large random graphs. Here we generalize this work and introduce a new algorithm which can be used to identify signaling pathways in transduction PINs. A detailed discussion is given in the Supporting Information (SI). Minimizing equation (1) gives access to *connected* networks which include reliable edges and, at the same time, nodes which are significantly differentially expressed (see Fig 1). This cost function could easily be generalized to other type of interactions, e.g. gene-based information such as results of knock-out experiments. Biological priors such as the relative position of proteins (nodes) along the tree or their expected degree could also be easily included in the same scheme.

One main difficulty with an optimization over connected sub-graphs is that the connectivity condition is global rather than local, i.e. *can not* be verified by a set of simple local checks over the graph. This problem is dealt with here by switching to a richer description of the sub-graph that includes an extra variable for each graph node, which essentially denotes when (if ever) nodes would be visited by an algorithm that explores the sub-graph from a given starting "root" node. While such representations are in one-to-one correspondence to connected sub-graphs, the connectivity condition *does* have an expression as a set of simple local conditions for the new variables.

The proposed algorithm consists in a set of functional equations for estimating the probabilities that individual links belong to the optimal sub-graph at a given distance from the starting node. Such equations can be written in a computationally efficient form that is solved by iteration in a so–called message–passing procedure. The general derivation and other details are given in SI text, whereas the source code is available for download [20]. The proof that in certain limit cases the algorithm provides optimal results can be found in [21].

Tests and data analysis. In order to assess the general efficacy of the algorithm, it was tested against the collection of MSTT benchmark problems in the SteinLib data set [22] which defines the state of the art in the field. Though the benchmarks problems are not of biological nature, they are both large and difficult to solve, and therefore particularly useful for comparing performance of different algorithms. Quite surprisingly the best known cost of almost all of the open problems could be improved in a fraction of the computational time (details and complete Tables in SI text). Most of the heuristic algorithms with the previously best known performance are based on Linear Programming (LP) relaxations complemented with preprocessing of the underlying graph and a branch-and-bound strategy, e.g. [15] (details also in SI). Further direct comparisons between such methods and our approach suggest that already for moderate size networks the LP-based methods become highly inefficient (details in SI).

As a second preliminary test, on biological data, we have compared our technique to another method for the inference of linear signaling pathways based on color-coding [11], the optimization criterion of which is a restriction of ours to the edge cost part. We have assessed our algorithm performance relative to [11] with the same data and optimization criterion. Essentially, the same pathways were found much more efficiently (due to the fact that the computational cost does grow only linearly with the chain length and not exponentially as in [11]), and it was possible to recover their variability by adjusting the chain length (details in SI text), thereby proving the capacity of our algorithm to recover known biology.

Pheromone response data. Finally, we have applied the algorithm to analyze pheromone sensing on a yeast protein network built by fusion of the MIPS[23] and DIP[24] networks, using 56 large-scale expression datasets created to reconstruct the pheromone pathway experimentally by studying the expression of strains deleted for key genes in the pathway [25]. This system was chosen as a case study in virtue of a pre-existing good theoretical understanding of its functioning. The pheromone response system is in fact a well-studied MAPK kinase cascade, which permits communication previous to mating in haploid yeasts. Upon sensing a pheromone, cell cycle is arrested, cytoskeleton and membrane structure are modified, and finally mating occurs by fusion of two cells of opposite sexual type.

The identification of optimal sub-networks was done for a range of λ values, giving us variable structures going from the backbone of the network to a very detailed picture of each sub-pathway. The root of the trees was defined to be STE2, the α-pheromone receptor and entrance of the entire pheromone pathway [26, 27].

The biological coherence of our results was assessed by showing that the average number of GO Slim annotations shared between neighbors in the inferred trees is significantly superior to those of random trees (t-test, $p < 5.10^{-7}$, see SI). The results obtained from each one of the 56 expression patters have been merged, leading to a final network with links which are weighted by their frequency of appearance. For each identified tree there is also a direction associated to each link toward the chosen root (not displayed in Fig. 3 for clarity).

The network we found (see Fig 3) contains known pathways, whose completeness depends on λ. It shows signaling going from STE2, to the cell cycle module (starting from CDC28), proteins involved in cell polarity and cytoskeleton reorganization (RVS161 and RVS167), and another branch containing essentially PRM proteins and proteins involved in sphingolipid synthesis (ELO1) or cell wall (CHS1). A closer look at the structure of the inferred sub-network shows that it is constituted of two types of proteins: proteins differentially expressed during pheromone response, already discovered by transcriptomic studies, and protein bridging between sub-parts of the network, while having a stable expression level in these conditions. These proteins, that will be called herein *Steiner proteins* in reference to Steiner nodes in the MSTT, are not detected by classical transcriptomics, but may have an informative or signaling role nonetheless: they allow signal propagation between modules of the transduction network, and can be discovered only through a combined analysis of the protein interaction and transcriptomics data.

One such protein we found is COS8, a gene of unknown function, which appear at the head of the third main branch starting from STE2, linking between the standard pheromone pathway and a group of proteins related to membrane structure. To ensure that COS8 frequent occurrence is not due to a bias of our optimization criterion to predilect proteins having high connectivity[12], we assessed the statistical significance of the Steiner proteins by bootstrap experiments with noisy costs and prizes. Interestingly, the COS8 protein appears frequently enough to be considered as significant and biologically relevant. Given the interest in uncovering a role for a member of a large gene family in yeast (the COS family contains 11 highly similar members), we therefore attempt to infer its biological role using its neighborhood in the inferred trees.

COS8 is suggested to be the target of the STE2-GPA1-SST2-SHO1 cascade. As SHO1 is also a sensor involved in the filamentous growth and osmolarity pathway, COS8 could be an actor of more than just the pheromone system, and could

provide cross-talks between these different pathways. Out of 48 proteins experimentally shown to interact with it, a small subset appear in most simulations, composed of the membrane protein PRM10, the fatty acid elongases ELO1, SUR4 and FEN1 – involved in the first steps of sphingolipid synthesis–, and other components of the secretion pathway and ceramide synthesis such a AUR1, LAC1 and IFA38. Such homogeneity in the uncovered interactions of COS8, relative to the diversity in the 48 interactants set, allows us to putatively predict both localization and role of COS8: it should be an endoplasmic reticulum (ER) protein involved in sphingolipid synthesis. Concerning localization, these predictions precise the results of a previous study using GFP fluorescence, that localized COS8 either in the ER or the nucleus[28]. Moreover, the single hint about COS8 function, uncovered in a large-scale transcriptomic experiment, and not experimentally verified, is an undetermined role in the unfolded protein response, a process occurring in the ER[29], in agreement with our results.

As for COS8 role, sphingolipids are an essential components of the membrane, being part of the lipid rafts microdomains [30]. To check the relevance of COS8 for the sphingolipid synthesis, and more generally for membrane structure, we analyzed with the same algorithm another large-scale dataset [31] testing for response to caffeine or rapamycin stress, which are also known to inhibit TOR pathway [32] and therefore to disrupt membrane structure. Consistently with the pheromone results, COS8 was also detected as a Steiner protein, and was significantly enriched in the resulting sub-graphs, with the same interacting partners. We therefore decided to check experimentally our predictions about COS8 interactions and putative role.
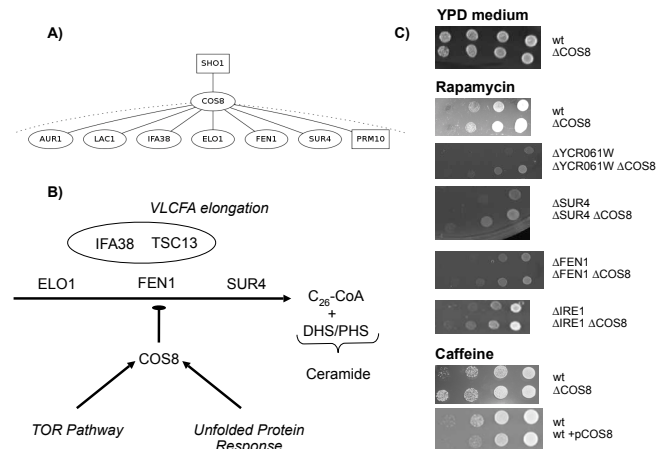
Fig. 2. A) The main proteins interacting with COS8 in our signaling tree. Squares stand for membrane proteins. The dotted lines show that many protein interactions are never found in our tree. B) A global scheme of the putative negative regulatory role of COS8, at the interface between very long chain fatty acid elongation (VLFCA), TOR signaling and the unfolded protein response (See text). C) A subset of the genetic experiments showing i) the rescue of ΔYCR061W, ΔSUR4, ΔFEN1, and ΔIRE1 by ΔCOS8 in rapamycin medium; ii) the effects of either the disruption or the over-expression of COS8 in a caffein medium

Strain	control	0.1 mM myriocin
wt	20	1.9
ΔCOS8	20.8	6.7

Experimental validation. We proceeded to genetic experiments in strains with disrupted COS8 or containing a plasmid over-expressing (Fig 2). We tested interactions with the main components of sphingolipid synthesis, in various conditions (full experimental details in SI text).

We made the ΔCOS8 strain for this study, by replacing the chromosomal copy of this gene by cos8::LEU2 deletion cassette. For the overexpression study, we cloned COS8 gene in multicopy plasmid pRS426(with URA3 marker). To obtain double mutants used in this study cos8::LEU2 deleted BY4741 strain was crossed with a strain from YKO collection (ORF::KAN). Resulted diploid was dissected and segregants having KAN and LEU2 markers were selected and used for further experiments. Details concerning all the experiments are given in SI.

We found that deleting COS8 rescue completely the phenotype of ΔFEN1 and ΔYCR061W strains in rapamycin medium – as well as caffein –, and partially the phenotype of ΔSUR4 in the same conditions. Oppositely, a strain over-expressing COS8 is hypersensitive to caffein and presents the same phenotype as FEN1 deletion. This shows a clear interaction between TOR signaling, COS8 and long chain fatty acid elongation, and hints at a negative regulation of very long chain fatty acid (VLCFA) elongation by COS8. As a control experiment, we checked for genetic interactions with IRE1, the master regulator of unfolded protein response, which is the annotated role of COS8[29], and we also found this interaction (See Fig. 2). Finally, we tested for growth defect in ΔCOS8 strains in media containing myriocin. Myriocin is a very potent inhibitor of serine palmitoyltransferase [33], the first step in sphingosine biosynthesis, and slow sphingolipid synthesis. We compared the sensitivity of ΔCOS8 and wt strains to this antibiotic. After 36 hours of incubation in liquid YPD media with 0.1 mM of myriocin we observed a transient but clear positive effect of COS8 deletion on cell growth (See Table).

Weak but reproducible increase of myriocin resistance for ΔCOS8 could also be seen on solid YPD media. This can be interpreted if COS8 indeed regulates negatively the sphingolipid production, via VLCFA synthesis : in these conditions the cell growth rate should become dependent on the efficiency of VLCFA elongation, which is less restricted in ΔCOS8 strains, therefore leading to smaller effects of myriocin. We therefore conclude that the function of COS8 is indeed related to sphingolipid metabolism, and probably regulates negatively the VLCFA synthesis and therefore the sphingolipid production.

Discussion

Algorithmic predictions and genetic experiments show interactions between sphingolipid synthesis, particularly ceramide, pheromone response and TOR signaling. Sphingolipids have recently been involved in the TOR-regulated network[34, 35]: TOR is able to activate the reaction of synthesis of ceramide from dihydrosphingosine. The molecular mechanisms of this regulation are still unknown, but are coherent with our experimental results about a negative role of COS8 in VLCFA elongation, as ΔCOS8 strains are resistant to caffein and ra-

pamycin – two components known to inhibit the TOR pathway. As sphingolipids are now discovered to be both essential membrane components and signaling molecules, understanding the regulation of their synthesis by various pathways, and the potential cross-talk they could provide, is a crucial issue. Here, regulation of sphingolipid synthesis by COS8 would provide the cell with the ability to integrate signal from the pheromone pathway, the osmolarity pathway, and the TOR pathway in order to modify its membrane structure. Finally, COS8 is a member of a gene family of 11 highly similar members. Further investigations would be needed to identify the functional role of other members of the family, considering that both our predictions and experiments seem to indicate that COS8 has a major effect, among all members of the COS family.

Conclusions

We have presented a new computational technique, inspired from statistical physics, which can efficiently extract new useful information about interactions in signaling pathways (from gene expression and protein-protein data) by solving an appropriately defined optimization problem on graphs. Our method not only provides candidate networks linking proteins of known function; the method also suggests new roles for proteins of previously unknown function. As a test case we specifically predict a functional role for the COS8 protein (a member of a large gene family with yet unknown functional role) both in sphingolipid synthesis and in the TOR pathway in *S. cerevisiae*. We have validated the prediction by providing experimental evidence showing that COS8 is involved in a regulatory loop at the level of ceramide synthesis. Our new algorithmic technique has several properties which should make it of significant value to optimization problems in systems biology: efficiency (nearly linear time complexity), simplicity (it is based on a fixed point equation), parallelizability, and the ability to include other biological priors, such as synthetic lethal interactions or phosphoproteomics data. Moreover, the technique outperforms the best-known techniques in the field: we tested it on unsolved instances of the best-known library (SteinLib), and it achieved better optima than were known previously. Finally, it is relatively easy to adapt the technique to a large class of network reconstruction problems, including many which arise in systems biology.

Materials and Methods

Full methods are in SI Text. They include: (1) Algorithm design (the model, derivation of the message-passing cavity equations, the max-sum limit, computation of marginals, iterative dynamics and reinforcement, a note on directness), (2) Numerical results on benchmark problems and direct comparison with LP-based techniques, (3) Data source and results (including data concerning GO annotation enrichment), (4) Experimental protocols (Strains, media and culture conditions, Construction of multi-copy plasmid with COS8 chromosomal allele, Construction of the COS8 deleted strain, Construction of double mutants, Drug sensitivity assays), (5) Algorithm comparison with previous data.

1. T.C. Elston (2008) Probing pathways periodically. Sci Signal 1, pe47.

2. A.J. King, D. R. Patrick, R. S. Batorsky, M. L. Ho, H. T. Do, S. Y. Zhang, R. Kumar, D. W. Rusnak, A. K. Takle, D. M. Wilson, E. Hugger, L. Wang, et al., (2006) Demonstration of a genetic therapeutic index for tumors expressing oncogenic BRAF by the kinase inhibitor SB-590885. Cancer Res 66, 11100.

3. J.-J. Pei and J. Hugon (2008) mTOR-dependent signalling in Alzheimer's disease. J Cell Mol Med 12, 2525.

4. M. Thattai, and A. van Oudenaarden (2002) Attenuation of noise in ultrasensitive signaling cascades. Biophys J 82, 2943.

5. B.N. Kholodenko (2006) Effects of sequestration on signal transduction cascades. Nat Rev Mol Cell Biol 7, 165.

6. J.W. Locasale and A. K. Chakraborty (2008) Regulation of signal duration and the statistical dynamics of kinase activation by scaffold proteins. PLoS Comput Biol 4, e1000099.

7. C.J. Bashor, N. C. Helman, S. Yan, and W. A. Lim (2008) Using engineered scaffold interactions to reshape MAP kinase pathway signaling dynamics. Science 319, 1539.

8. B.A. Benayoun, and R. A. Veitia (2009) A post-translational modification code for transcription factors: sorting through a sea of signals. Trends Cell Biol 19, 189.

9. B. Soufi, C. D. Kelstrup, G. Stoehr, F. Fröihlich, T. C. Walther, and J. V. Olsen (2009) Global analysis of the yeast osmotic stress response by quantitative proteomics. Mol Biosyst 5, 1337-1346.

10. M.S. Scott, T. Perkins, S. Bunnell, F. Pepin, D. Y. Thomas, and M. Hallett (2005) Identifying regulatory subnetworks for a set of genes. Mol Cell Proteomics 4, 683.

11. J. Scott, T. Ideker, R. M. Karp, and R. Sharan (2006) Efficient algorithms for detecting signaling pathways in protein interaction networks. J Comput Biol 13, 133.

12. A. White, and A. Mayan (2008) Connecting seed lists of mammalian proteins using Steiner trees . Nature Precedings.

13. X.-M. Zhao, R.-S. Wang, L. Chen, and K. Aihara (2008) Uncovering signal transduction networks from high-throughput data by integer linear programming. Nucleic Acids Res 36, e48.

14. M. Bailly-Bechet, A. Braunstein and R. Zecchina (2009) A prize-collecting Steiner tree approach for transduction network inference. P. Degano and R. Gorrieri (Eds.): CMSB 2009, LNBI 5688.

15. I. Ljubic, R. Weiskircher, U. Pferschy, G. Klau, P. Mutzel, and M. Fischetti (2006) An algorithmic framework for the exact solution of the prize-collecting Steiner tree problem. Mathematical Programming, B 105, 427.

16. M. Dittrich, G. Klau, A. Rosenwald, T. Dandekar, and T. Muller (2008) Identifying functional modules in protein-protein interaction networks: an integrated exact approach. Bioinformatics 24, 223.

17. S.-S.C. Huang, and E. Fraenkel (2009) Integrating proteomic, transcriptional, and interactome data reveals hidden components of signaling and regulatory networks. Sci Signal 2, ra40.

18. N. Yosef, L.Ungar, E. Zalckvar, A. Kimchi, M. Kupiec, E.Ruppin, and R. Sharan (2009) Toward accurate reconstruction of functional protein networks. Molecular Systems Biology 5, 248.

19. M. Bayati, C. Borgs, A. Braunstein, J. Chayes, A. Ramezanpour, and R. Zecchina (2008) Statistical Mechanics of Steiner Trees. Phys Rev Lett 101, 037208.

20. www.polito.it/cmp

21. M. Bayati, A. Braunstein, and R. Zecchina (2008) A rigorous analysis of the cavity equations for the minimum spanning tree. J. Math. Phys. 49, 125206.

22. T. Koch, A. Martin, and S. Voss (2000), SteinLib: An Updated Library on Steiner Tree Problems in Graphs, Technical Report ZIB-Report 00-37, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Takustr. 7, Berlin.

23. U Güldener, M Münsterkötter, M Oesterheld, P Pagel, A Ruepp, H-W Mewes, and Stümpflen (2006) Mpact: the mips protein interaction resource on yeast.Nucleic Acids Res, 34 (Database issue), D436.

24. I Xenarios, D W Rice, L Salwinski, M K Baron, E M Marcotte, and D Eisenberg (2000) Dip: the database of interacting proteins. Nucleic Acids Res, 28(1), 289.

25. C.J. Roberts, B. Nelson, M. J. Marton, R. Stoughton, M. R. Meyer, H. A. Bennett, Y. D. He, H. Dai, W. L. Walker, T. R. Hughes, M. Tyers, C. Boone, et al. (2000) Signaling and circuitry of multiple MAPK pathways revealed by a matrix of global gene expression profiles. Science 287, 873.

26. D. D. Jenness, A. C. Burkholder, and L. H. Hartwell (1983) Binding of alpha-factor pheromone to yeast a cells: chemical and genetic evidence for an alpha-factor receptor. Cell 35, 521–529.

27. A. C. Burkholder and L. H. Hartwell (1985) The yeast alpha-factor receptor: structural properties deduced from the sequence of the STE2 gene. Nucl Acids Res 13, 8463–8475.

28. I. Spode, D. Maiwald, C. P. Hollenberg, and M. Suckow (2002) ATF/CREB sites present in sub-telomeric regions of Saccharomyces cerevisiae chromosomes are part of promoters and act as UAS/URS of highly conserved COS genes. J Mol Biol 319, 407.

29. K.J. Travers, C. K. Patil, L. Wodicka, D. J. Lockhart, J. S. Weissman, and P. Walter (2000) Functional and genomic analyses reveal an essential coordination between the unfolded protein response and ER-associated degradation. Cell 101, 249.

30. R.C. Dickson, and R. L. Lester (2002) Sphingolipid functions in Saccharomyces cerevisiae. Biochim Biophys Acta 1583, 13.

31. K. Kuranda, V. Leberre, S. Sokol, G. Palamarczyk, and J. François (2006) Functional and genomic analyses reveal an essential coordination between the unfolded protein response and ER-associated degradation. Mol Microbiol 61, 1147.

32. X.F. Zheng, D. Florentino, J. Chen, G. R. Crabtree, and S. L. Schreiber (1995) TOR kinase domains are required for two distinct functions, only one of which is inhibited by rapamycin. Cell 82, 121.

33. Y. Miyake, Y. Kozutsumi, S. Nakamura, T. Fujita, and T. Kawasaki (1995) Serine palmitoyltransferase is the primary target of a sphingosine-like immunosuppressant, ISP-1/myriocin. Biochem Biophys Res Commun 211, 396.

34. S. Aronova, K. Wedaman, P. A. Aronov, K. Fontes, K. Ramos, B. D. Hammock, and T. Powers (2008) Regulation of ceramide biosynthesis by TOR complex 2. Cell Metab 7, 148.

35. C.J. Mousley, K. Tyeryar, K. E. Ile, G. Schaaf, R. L. Brost, C. Boone, X. Guan, M. R. Wenk, and V. A. Bankaitis (2008) Trans-Golgi network and endosome dynamics connect ceramide homeostasis with regulation of the unfolded protein response and TOR signaling in yeast. Mol Biol Cell 19, 4785.
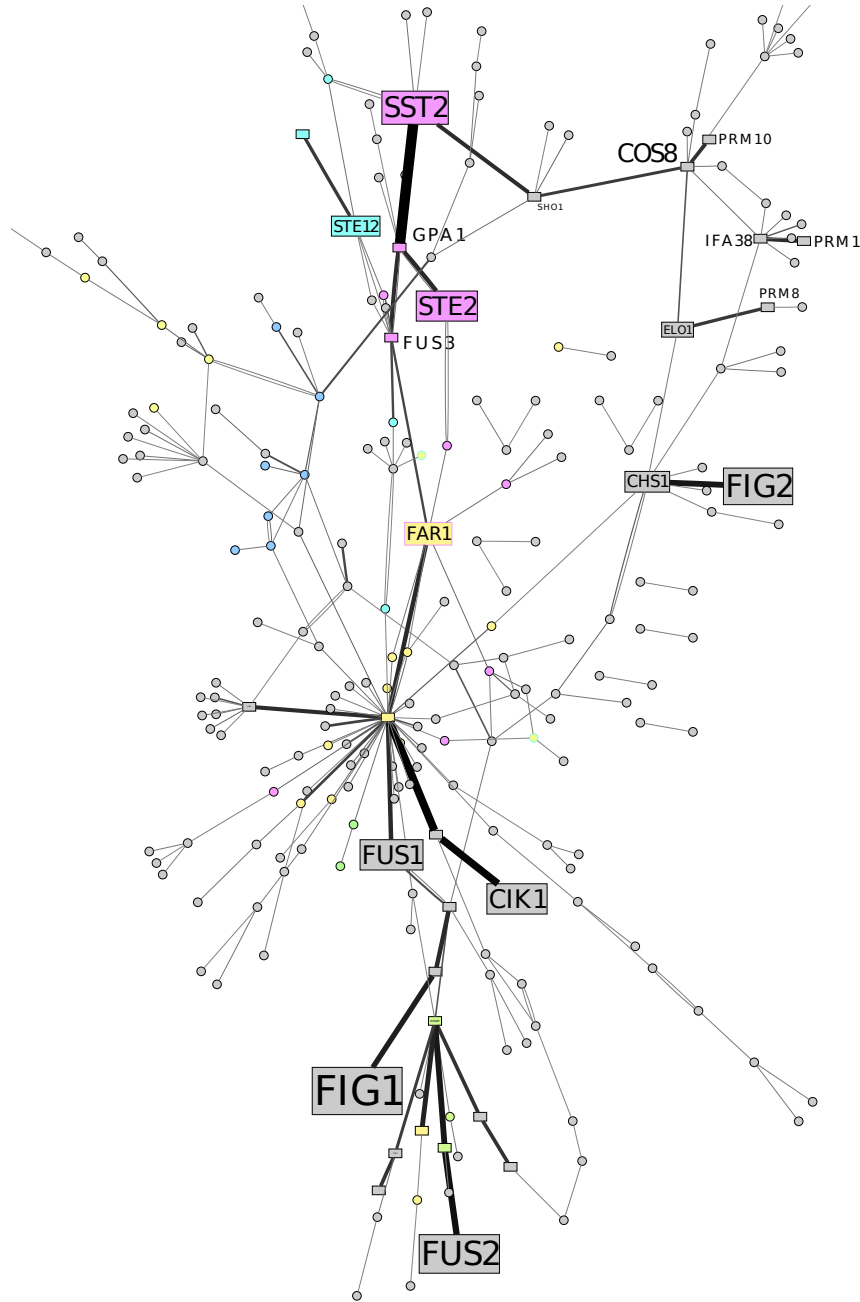
Fig. 3. This graph is a sub-network of the protein-protein interaction map, obtained by including nodes that appear more than 30% of the times on the 56 inferred Steiner trees for $\lambda = 0.2$, with link intensity proportional to the number of times the specific connection was found, and node size is proportional to average prize. The layout was decided in order to minimize crossings with the Graphviz suite. Afterward, colors were added denoting the main GO annotation: Actin (light green), Cell Cycle (yellow), Chromatin structure (blue), Spindle Checkpoint (dark green), Cell wall (cyan) and Pheromone sensing (magenta). The annotations were obtained from the SGD project "Saccharomyces Genome Database".